CP3-Origins-2014-001 DNRF90
DIAS-2014-1

# BSMBench: a flexible and scalable supercomputer benchmark from computational particle physics

Ed Bennett[1a], Luigi Del Debbio[2b], Kirk Jordan[3c], Biagio Lucini[4a], Agostino Patella[5d,e], Claudio Pica[6f], and Antonio Rago[7e]

[a] *Department of Physics, Swansea University, Singleton Park, Swansea, SA2 8PP UK*
[b] *The Higgs Centre for Theoretical Physics, University of Edinburgh, Edinburgh, EH9 3JZ UK*
[c] *IBM Research, Thomas J Watson Research Center, 1 Rogers Street, Cambridge, MA 02142 USA*
[d] *PH-TH, CERN, CH-1211 Geneva 23, Switzerland*
[e] *School of Computing and Mathematics & Centre for Mathematical Science, Plymouth University, Drake Circus, Plymouth, PL4 8AA UK*
[f] *CP-Origins & the Danish IAS, University of Southern Denmark, Odense, Denmark*

arXiv:1401.3733v1 [cs.DC] 15 Jan 2014

January 16, 2014

## Abstract

Benchmarking plays a central role in the evaluation of High Performance Computing architectures. Several benchmarks have been designed that allow users to stress various components of supercomputers. In order for the figures they provide to be useful, benchmarks need to be representative of the most common real-world scenarios. In this work, we introduce BSMBench, a benchmarking suite derived from Monte Carlo code used in computational particle physics. The advantage of this suite (which can be freely downloaded from `http://www.bsmbench.org`) over others is the capacity to vary the relative importance of computation and communication. This enables the tests to simulate various practical situations. To showcase BSMBench, we perform a wide range of tests on various architectures, from desktop computers to state-of-the-art supercomputers, and discuss the corresponding results. Possible future directions of development of the benchmark are also outlined.

## 1 Introduction

Evaluating the performance of different architectures is an important aspect of High-Performance Computing (HPC). This evaluation is most often performed using benchmark applications, which conduct simulated workloads of known complexity, so their execution times may be compared between machines. Historically, the High-Performance LINPACK (HPL) [1] benchmark has been the "gold standard" of HPC benchmarks, with the Top500 Supercomputing Sites [2] (listing the machines with the top HPL scores) being widely used as a proxy for the "fastest" computers in the world.

Despite remaining the universal standard for HPC benchmarking, HPL has some shortcomings that are starting to cast serious doubts about its current and future usefulness to assess the real capability of a given machine for realistic workloads. From an algorithmic point of view, HPL is a rather simple benchmark. This has led to machines' architectures (and compilers) being tweaked to optimise their performance in the HPL benchmark. Until recently this has not been a problem, with workloads' requirements correlating roughly with those of HPL. However, as highlighted in a recent paper by Dongarra and Heroux [3], as we look towards machines capable of tackling exascale problems, the demands of workloads have begun to diverge from those of HPL, to the point that design decisions made to improve HPL performance is detrimental to the machines' performance in real-world tasks.

Additionally, the time taken to complete a run of HPL has increased with machine size, with current top-end machines requiring over 24 hours; this is expensive, in terms of both highly-constrained machine time and energy usage, leading to single runs being performed, rather than multiple runs to confirm consistency as might be made with a shorter-running tool. Dongarra and Heroux propose a new HPCG benchmark based on Conjugate Gradient (CG)

[1] pyedward@swan.ac.uk
[2] luigi.del.debbio@ed.ac.uk
[3] kjordan@us.ibm.com
[4] B.Lucini@swan.ac.uk
[5] Agostino.Patella@cern.ch
[6] pica@cp3.dias.sdu.dk
[7] Antonio.Rago@plymouth.ac.uk

matrix inversion with preconditioning to supplement HPL as a reference benchmark for more realistic workloads.[1]

In this paper, we propose a new benchmark based on code from Beyond the Standard Model (BSM) Lattice Gauge Theory (LGT), which fulfils many of the criteria laid out by Dongarra and Heroux for the HPCG benchmark. Algorithmically, our benchmark can be seen as a generalisation of performance tests derived from a branch of theoretical particle physics known as Quantum ChromoDynamics (QCD) on the Lattice (or Lattice QCD). For this family of scalable benchmarks, characterised by simultaneously high demands in terms of both computation and communication, there is a wide literature. The advantage of our benchmark over standard Lattice QCD-based ones is that it provides the additional benefit of being able to tune the ratio of communication over computational demands of the algorithm, by altering the physical theory that is being considered. This is an important feature that enables us to change the benchmark requirements in order to simulate various realistic workloads; no existing benchmarks based on Lattice QCD have this ability.

The rest of the paper is organised as follows. To place our results in context, in Section 2, a brief account of the approach to benchmarking underlying our work is given. In Section 3, an overview of the physical context from which the code has originated is provided. Section 4 describes the structure of the benchmark code and the benchmarking strategy it employs. Section 5 shows the results we have found when running the benchmark on a selection of systems, including IBM Blue Gene machines, x86-based clusters, and a desktop workstation. In Section 6 we discuss some of the interesting features of the data found, and present suggestions for future work. Finally, Appendix A gives some technical details of the theoretical and computational calculations performed in BSM Lattice physics that are relevant to the benchmark.

## 2 Benchmark requirements

Before discussing our benchmark and the physical problem from which it was spawned, it is worth setting the scene by specifying our requirements for the benchmarking tool and—as a consequence—the scope of the results discussed in this work. Various components of a supercomputer (or, in fact, of any computer) can be benchmarked, either (at least ideally) in isolation, or in reciprocal interaction. In our case, we are interested in the speed of execution of a complex numerical problem (i.e. a problem requiring a high number of floating point operations) in a scenario in which the graphic rendering does not play any role and disk access times are negligible. Hence, the main parameter we will focus on is the number of operations that the CPU can execute, which is influenced both by the speed of the CPU and the speed of access to memory. In addition, since the problem can be distributed across interlinked CPUs that do not necessarily share the same physical memory, speed of communication of data from one computational unit to another and latencies of interconnections are also important. Applications with those requirements are commonly found in physics, engineering and finance, to mention only a few disciplines in which High Performance Computing plays a key role, and the results we shall discuss are relevant in assessing the performance of a supercomputing architecture for this class of problems. On the other hand, by design our benchmark will not capture the performance of a system for applications that are—for instance—disk-intensive, nor can it help in assessing the energy efficiency of a machine.

A benchmark that is able to assess the performance (as defined above) of an HPC system should also fulfil additional criteria of reliability, practicality, and generality, among which are the following:

- It should be highly portable.
- It should give the correct mathematical answer on the tested architecture.
- It should use code that is representative of the main body of the computation being benchmarked.
- It should give output that can be read by people who are not knowledgeable in the area of the benchmark, so that system builders can run the benchmark internally and understand its output.
- It should be sufficiently large that the machine in question can demonstrate its power (i.e. not so short that initialisation times dominate, or so small that only a minor portion of a machine can be used).

[1] Other supplementary benchmarks have been introduced in recent years to assess performance at graph-traversing applications [4] and energy efficiency [5]; however, neither of these is targeted at realistic compute-intensive workloads, and they will not be discussed in this work.

- It should run in reasonable time (30 minutes – 1 hour is ideal, although due to its age and design LINPACK can take over 24 hours to run on the largest current machines; there is active research into how to reduce this time whilst still maintaining a comparable output).
- It should be reliable and repeatable (if one machine is rated as having better performance than another, it should always do so for that metric, and repeated runnings of a test should give the same score).

Following these main guiding principles, the benchmark discussed in this paper, BSMBench, has been designed and developed to be a robust, reliable, scalable and flexible HPC performance assessment tool that can be used to gather information for a wide variety of systems in a wide range of appropriate workloads.

# 3 Lattice Gauge Theory and HPC benchmarking

In order to understand the properties of the algorithm we are proposing and how they can be used in benchmarking, in this section we give a very gentle and short introduction to Lattice Gauge Theory, and in particular to the physical problems it intends to address and the numerical methods it implements to solve it. We refer to the appendix and to the quoted literature for the more technical details.

## 3.1 QCD and novel strong forces

Lattice Gauge Theory is a framework originally devised to solve QCD. QCD is the theory describing physics below the level of atomic nuclei. All nuclear matter (i.e. protons and neutrons) is made of quarks and gluons, with the gluons transmitting the (strong nuclear) force that binds the quarks together. The strong force is one of the four fundamental forces of nature, along with the weak force, electromagnetism (which taken together comprise the Standard Model of particle physics) and gravity, and is also responsible for the interactions of nuclear matter.

While the Standard Model is a remarkably successful theory, there remain a few questions it leaves unanswered. It is known that the electromagnetic and weak interactions are unified at high energies to a single "electroweak" interaction. The simplest explanation of how this breaks down to the two forces we experience at low energies is the existence of an Higgs field, which also gives rise to a scalar boson. Such a state was observed at the Large Hadron Collider at CERN in 2012 [6, 7]. However, the Higgs theory also introduces a problem, in that the mass of the Higgs boson is "unnaturally" light. Theorists therefore are looking "beyond" the Standard Model, to see whether any alternative explanations may be found. One such avenue of research suggests that rather than a Higgs-like theory, a QCD-like theory is responsible for the electroweak symmetry breaking.

## 3.2 A first look at the lattice

Lattice gauge theory is the set of techniques that allow QCD and other BSM theories to be simulated computationally. A recent pedagogical introduction to Lattice gauge theory (mainly focused on QCD) is provided in [8]. As is common in scientific problems, the continuous space-time of QCD is replaced by a discrete four-dimensional *lattice* of points, with continuous derivatives replaced by finite differences. Since QCD is defined in terms of a finite four-dimensional space-time (three spatial dimensions, plus one time), the lattice is also four-dimensional.

Fields are functions of the position, and may live on the sites or the links between them. Their values may be real or complex numbers, or more complicated objects—in the case of QCD and BSM theories of interest, the objects are complex $N \times N$ matrices with determinant equal to one. These matrices form a group, which is called the special unitary $N \times N$ group, referred to as SU($N$). The parameter $N$ is an integer number specifying the physical theory. In this context, QCD is the theory described by SU(3). Quark fields, living on the lattice sites, require a different form of SU($N$) matrix to the gluon fields, which live on the links, and can have various dimensions, corresponding to the representation of SU($N$) they transform according to (the fundamental representation of SU(3) in the case of QCD; other representations we shall consider in the more general case are the adjoint, the symmetric and the antisymmetric representations). They come in $N_{\mathrm{f}}$ copies (*flavours*) with the same mass or different masses. A field's value at a position is related (in a non-trivial manner) to the probability of observing a particle at that point.

The set of values of all fields across the lattice is known as a *(field/gauge) configuration*, and

can be seen as the value of the variables of the theory at a given time (i.e. as a snapshot of the system for some fixed time). The evolution of the matter and field variables is determined by the dynamics derived from the physical theory according to well known laws; the numerical simulation realises this evolution[2]. Monte Carlo methods (including but not limited to Metropolis and Molecular Dynamics algorithms) are used to generate new gauge configurations from old, such that the set of configurations has an appropriate statistical distribution of values. With a suitably large number of configurations, physical quantities of interest (*observables*) may be calculated using appropriate programs. Since not all observables that can be measured in reality are easily calculable on the lattice, much research focuses on devising new methods to calculate observables efficiently.

### 3.3 Computations

The gauge and fermionic degrees of freedom (which in QCD are the gluon and quark fields respectively) are stored in structured arrays. In particular, the fermion variables can be integrated out. After this has been performed, the computational kernel of the problem consists of the inversion on a particular vector of a sparse matrix $D$ (referred to as the *Dirac operator*) whose entries are SU($N$) matrices. In order to find the wanted inverse, various algorithms are available (e.g. Conjugated Gradient, Minimum Residues, Lanczos, etc.), which mostly consists in a recursive procedure [9]. This recurrence can be accelerated using techniques such as domain decomposition and preconditioning [9].

Being local, the problem is well suited for parallelisation [8]. For portability reasons, most often this parallelisation uses the Message Passing Interface, or MPI in short (see the next section), dividing the lattice into a coarse grid of space-time sub-lattices. Each sub-lattice includes a boundary layer at least one lattice site thick in each direction, to allow computation of derivatives without needing to fetch data from other processes. These must however be updated regularly as the values are changed by adjoining processes. This is part of the communications demand of the code; the other is that the update processes use quantities which depend on fields at all lattice sites, rather than just the local sub-lattice. This aspect is minimised, however, by computing these quantities piecewise locally, and collecting the pieces from each process. Nevertheless, the computation demands are such that not only is the bandwidth important, but also network latencies play a major role.

We begin to see now that adjusting the $N$ of SU($N$) and the fermion representation will each alter the compute and communications demands of the task.

## 4 BSMBench

### 4.1 The underlying physics code

BSMBench is based on a Lattice Gauge Theory research code called HiRep, which has been developed over the years by some of the authors of this paper. Compared to standard Lattice QCD codes[3], in which $N = 3$ is often hard-wired to simulate a theory of SU(3) matrices, HiRep has the ability of using $N$ as a parameter, which makes it a useful exploration tool for the physics of possible BSM theories. The original serial version of HiRep was firstly used in [12], where its operations are described in detail. With the growing computational demands of the research problem it intended to solve, a parallel version was developed that was first used to produce the results reported in [13–15]. Since then, much work has been published incorporating calculations produced with HiRep and code built upon it (for a non-exhaustive set, see for instance [16–21]).

The HiRep code comprises:

- a set of C libraries/modules containing shared functions;
- a set of C and Fortran programs for specific computations, making use of the above libraries;
- C++ code to generate header files for a particular fermion representation and gauge group at compile time;
- makefiles to link the above together;
- additional statistical analysis tools (C++ code with bash wrappers) to extract physical quantities from the program output.

Compile-time parameters (e.g. gauge group, fermion representation, boundary conditions,

[2] It is worth remarking that—in contrast to the more familiar behaviour in Classical Mechanics—in Quantum Mechanics and in Quantum Field Theory the evolution is not deterministic, but probabilistic.

[3] See for instance [10, 11] for examples of Lattice QCD codes.

use of MPI) are specified in a flags file. Runtime parameters (e.g. lattice volume and parallelisation, coupling and bare mass, directory for writing configurations, algorithmic parameters) are set in an input file, which takes the form of *key* = *value* pairs, one per line. A few parameters (e.g. run log file location) are set as command-line arguments.

As is common for lattice codes, parallelisation occurs using MPI, with the division splitting the lattice volume into chunks, one per process. In the absence of compiler-specific automatic multithreading optimisations, each MPI process is single-threaded. HiRep supports operating in single- and double-precision, but the precision of current research relies on double-precision arithmetic. BSMBench considers only double-precision floating-point variables. Other improvements currently being implemented in order to take advantage of more recent architectures include hybrid MPI-OpenMP programming and porting to GPUs and to Intel's Xeon Phi architecture.

## 4.2 The benchmark suite

The usefulness of using codes derived from Lattice QCD as generic benchmarking tools has been known for some time [22] and various benckmarks are available that use Lattice QCD techniques (see for instance [23]). In Lattice QCD, communication and computational demands are roughly balanced, which means that in a benchmark based on Lattice QCD a deficiency in one of the two areas could be masked by a good performance in the other. Instead, looking at different BSM theories changes the relative requirements of the ratio computation over communication. Because of the increased flexibility, a benchmark derived from BSM Lattice theories lends itself more to applications to areas outside of Lattice physics.

In order to obtain valuable information on both speed of computation and data transfer capabilities, we choose to make three different regimes available to benchmark:

- A communications-intensive regime, SU(2) with adjoint flavours, referred to as the "comms" test,
- A balanced regime, SU(3) with fundamental flavours, referred to as the "balance" test, and
- A compute-intensive regime, SU(6) with fundamental flavours, referred to as the "compute" test,

where in each case the value of $N$ and the fermion representation is fixed. More tests can be added starting from the HiRep code.

The strategy adopted for the benchmark is adapted from that of Lscher [23]. The program's preparation and execution occur as follows:

- Preparation:
  - The user prepares or selects a machine configuration file specifying their C compiler and appropriate flags; some such files are included with the distribution for machines with which BSMBench has been tested.
  - (Optional) The user may set a flag to indicate that they do not require a parallel version of the code.
  - The user calls the compile script with the parallel (and optionally the non-parallel) machine configuration file as argument(s).
  - The compilation script copies the machine configuration file into the appropriate place and sets appropriate flags in the HiRep configuration files.
  - The compilation script then calls the make command to generate the three components of the benchmark listed above, in parallel (unless told otherwise) and (optionally) non-parallel versions.
- Execution:
  - The user (or job control system) calls (via mpirun if necessary) the appropriate benchmark executable, pointing it at one of the supplied parameter sets.
  - The code reads in the parameter set, defining the geometry (i.e. the allocation of the lattice across the various processes) appropriately.
  - The code allocates a random gauge field and as many random fermion fields as are specified in the parameter set.
  - As a consistency check, the code inverts the Dirac operator[4] on the first random spinor[5] field and ensures

[4]The Dirac operator (briefly exposed in the Appendix) is the main computational kernel in Lattice Field Theory calculations involving fermions.

[5]A spinor is a mathematical object describing fermions. For the purposes of this paper, a spinor may be thought of as a vector.

that the residual drops below a given threshold. The input and output spinors are then re-randomised. This test is only carried out if requested in the parameter set, since it would dominate the execution time on smaller machines (e.g. workstations), where it will generally be unnecessary. The total number of iterations and overall time taken is recorded.

- The code calculates the square norm of each random spinor field in turn a given number of times, repeating until a time threshold is passed. The number of iterations starts at one, and doubles each time the process repeats, to avoid spending excessive time checking the progress.
- The code performs the operation $\psi_2 = \psi_2 + c\psi_1$, where $c$ is a complex scalar constant, on successive pairs of random spinor fields, repeating in the same manner as above.
- The code performs a Dirac operator application $\psi_2 = D\psi_1$ on successive pairs of random spinor fields, repeating in the same manner as above.
- The code uses its results and reference and output performance information. Performance is outputted in FLOP/s (Floating-Point OPerations per second) for numerical applications and for convenience the comparison with a reference machine (chosen to be one rack of BlueGene/P) is also provided.

For the purposes of comparing machines' performance, the Dirac operator application test is the most revealing.

The structure of the code is as follows. In the source distribution, there is:

- A directory containing a stripped-down HiRep distribution, with the addition of the benchmark code and comments on the licensing terms,
- A directory containing sample machine configuration files,
- A directory containing parameter sets for the default set of tests,
- An empty directory to hold output files,
- The README and LICENSE text files , and
- The compilation script.

The source code for BSMBench can be downloaded from `http://www.bsmbench.org`, where its development (open to contributions from any interested party) is hosted. The compilation script relies on a working MPI development environment able to compile C/C++ source code through the popular *make* build system. The compilations scripts have been tested on various Unix-like operating systems (among which IBM AIX, GNU/Linux, and OS X), and should work with minor variations on other systems. For simplicity, generic and architecture-specific template makefiles are also provided.

The supplied tests work with a four-dimensional lattice of size $64 \times 32^3$. This choice allows for comparison from workstation-class machines up to small supercomputers. Comparing results of tests at different lattice sizes is not recommended, since the associated overheads change. Additionally, the maximum parallelisation recommended is that giving a local lattice of $8 \times 4^3$; further parallelisation unreasonably increases the overheads. This limits the stock tests to 4096 MPI processes; equivalent to a single Blue Gene/P rack in Virtual Node mode. This limit could be removed by choosing a larger lattice size, as is done in the next section for the extended tests of Blue Gene/Q; however, this would make running the tests impractical on workstations. 4096 MPI processes is sufficient as a proof-of-concept and a test of moderately-sized machines.

# 5 Results

In order to prove its capabilities, the code has been used to systematically test a number of machines:

- Two clusters provided by HPC Wales, one Intel Westmere-based and the other powered by the Intel Sandy Bridge architecture;
- The BlueIce2 cluster at Swansea University (Intel Westmere-based);
- An IBM Blue Gene/P machine (which was used as the reference configuration); runs were split between the UKQCD machine in Swansea (part of the STFC DiRAC HPC facility) and a similar machine operated by IBM Research in Yorktown Heights, New York;
- An IBM Blue Gene/Q machine, also at IBM Research in Yorktown Heights;

| Machine | Processor | Number of cores | Threads per core | RAM | L2 cache | L3 cache | Interconnect | Compiler | Upper bound |
|---|---|---|---|---|---|---|---|---|---|
| HPC Wales (Westmere) | Intel Westmere Xeon X5650 | 12 | 2 | 36GB | 6MB | 24MB | InfiniBand | GCC 4.6.2 | Queue size |
| HPC Wales (Sandy Bridge) | Intel Sandy Bridge Xeon E5-2690 | 16 | 2 | 64GB | 4MB | 20MB | InfiniBand | Intel | Queue size |
| BlueIce2 | Intel Westmere Xeon X5645 | 12 | 2 | 24GB | 6MB | 24MB | InfiniBand | GCC 4.4.6 | Machine size |
| IBM Blue Gene/P | 32-bit PowerPC at 850MHz | 4 | 1 | 4GB | 8kB | 4MB | High-speed 3D torus | IBM XL C 9.0 for Blue Gene | Largest test size |
| IBM Blue Gene/Q | 64-bit PowerPC at 1600MHz | 16 | 4 | 16GB | 32MB | — | High-speed 5D torus | IBM XL C 12.0 for Blue Gene | Largest test size and machine size |
| ULGQCD | AMD Opteron 6128 | 8 | 1 | 96GB | 8MB | 24MB | Gigabit Ethernet | GCC 4.1.2 | Slow interconnects[a] |
| Mac Pro[b] | Intel Westmere Xeon X5650 | 12 | 2 | 24GB | 6MB | 24 MB | — | LLVM-GCC 4.2.1 | Machine size |

Table 1: The seven machines tested by BSMBench and details of their node configuration, divided into the groups in which they were plotted.

- The ULGQCD cluster at Liverpool University, part of the STFC DiRAC HPC facility;
- A Mac Pro, to include a high-level workstation in the comparison.

Relevant details of the machines are listed in table 1.

The primary test performed was running each of the benchmark tests on each machine at increasing partition sizes, from the minimum on which the tests would run (with node memory being the primary constraint) up to an upper limit determined by the factors listed in table 1. All tests were run with one MPI process per processor core—sometimes referred to as Virtual Node (VN) mode—since HiRep and BSMBench are not currently multithreaded. In standard implementations, most queue managers ignore any concurrent thread support of the CPUs, since MPI rarely gains an advantage by running multiple concurrent threads; multithreaded code must be written or compiled in to use this ability. Blue Gene/Q, however, does provide for this, thus up to 64 MPI processes may be run on a single processor.

Two other tests were run: firstly, a customised set of tests with a $128 \times 64^3$ lattice on Blue Gene/Q to allow a rack-for-rack comparison with a Blue Gene/P; this was again run in VN mode. Finally, a set of tests were run on a single Blue Gene/Q partition, varying the number of MPI processes per node from 1 to 64 in powers of 2, to test the performance of the multiple hardware threads of the Blue Gene/Q processor.

The plots are divided into groups by machine class to avoid overcrowding them. The primary battery of tests, shown in figures 1–8, shows a roughly linear relationship between FLOP/s performance and number of MPI processes (i.e. processor cores in use) in all cases. All machines show a slow drop in performance per node as the number of nodes is increased, associated with the increased overheads. The ULGQCD cluster shows a rapid tail-off in performance once the parallelisation exceeds the size of one node; for the largest size tested not only does the per-core performance drop, but the overall performance of the partition drops.

The rack-for-rack comparison, shown in figure 9 shows that for larger lattices Blue Gene/Q continues to scale as we expect. The comparison of different processor subdivisions, shown in figure 10 shows linear scaling up to 1 process per core. When further subdividing the cores, the

Key to symbols:

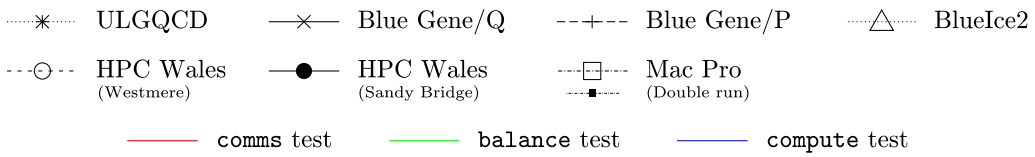


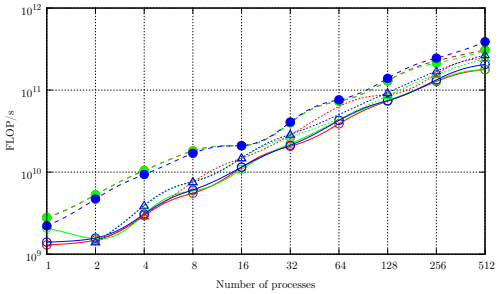


(a) HPC Wales and BlueIce2

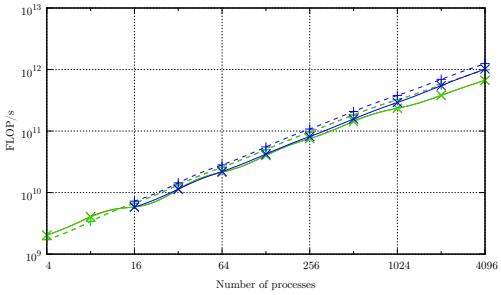


(b) Blue Gene

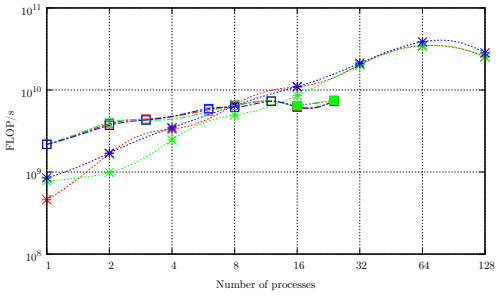


(c) Others

Figure 1: Results of the spinor field square norm test plotted for the whole ensemble.

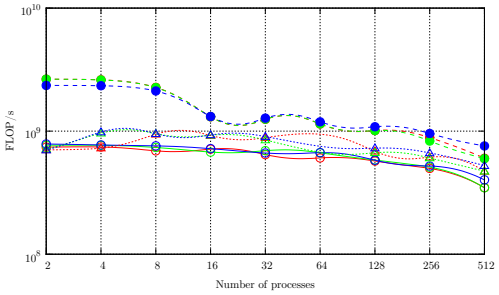


(a) HPC Wales and BlueIce2

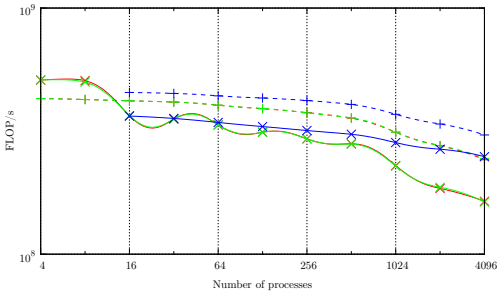


(b) Blue Gene

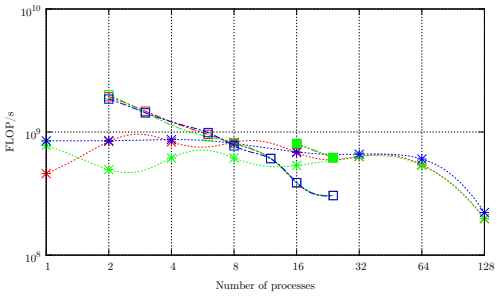


(c) Others

Figure 2: Results of the spinor field square norm test, plotting the average performance per MPI process.

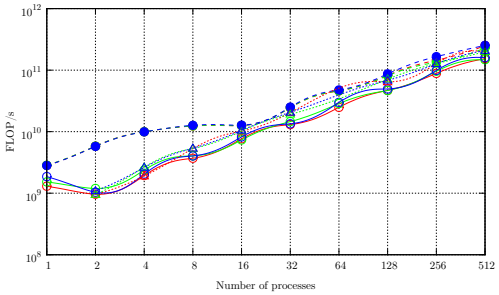


(a) HPC Wales and BlueIce2

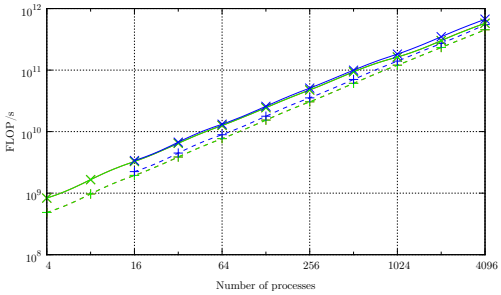


(b) Blue Gene

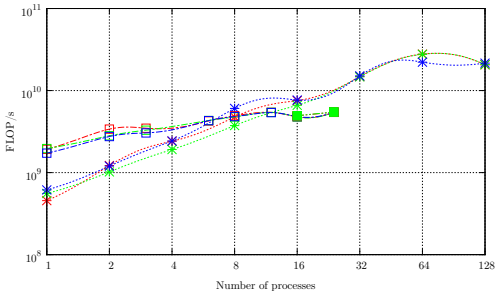


(c) Others

Figure 3: Results of the spinor field multiply–add, plotted for the whole ensemble. Symbols as on page 8.

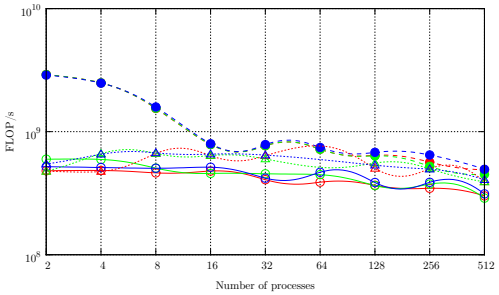


(a) HPC Wales and BlueIce2

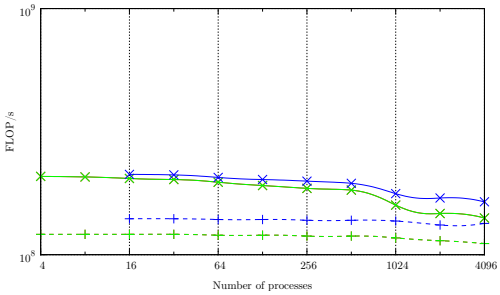


(b) Blue Gene

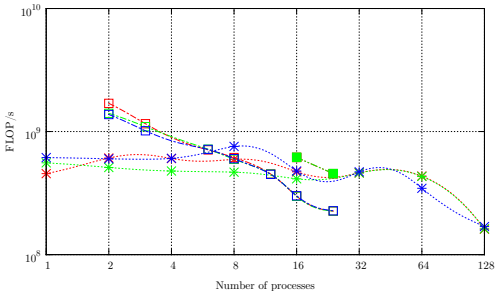


(c) Others

Figure 4: Results of the spinor field multiply–add, plotting the average performance per MPI process. Symbols as on page 8.

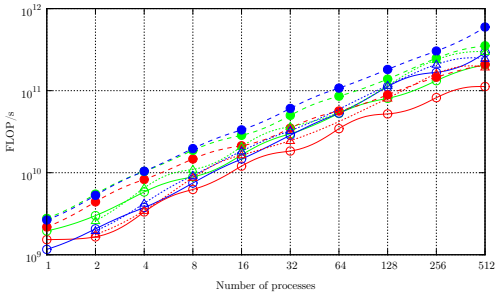


(a) HPC Wales and BlueIce2

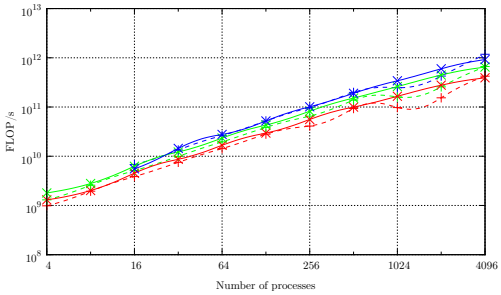


(b) Blue Gene

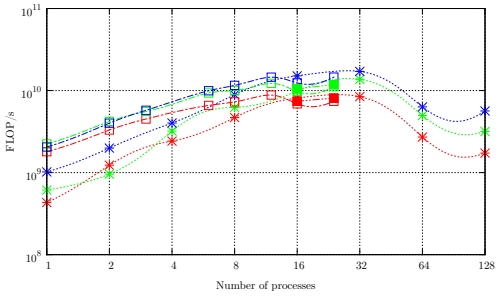


(c) Others

Figure 5: Results of the Dirac operator application test, plotted for the whole ensemble. Symbols as on page 8.

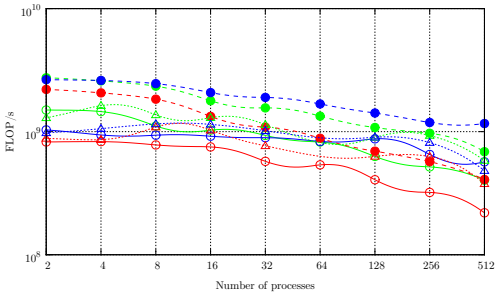


(a) HPC Wales and BlueIce2

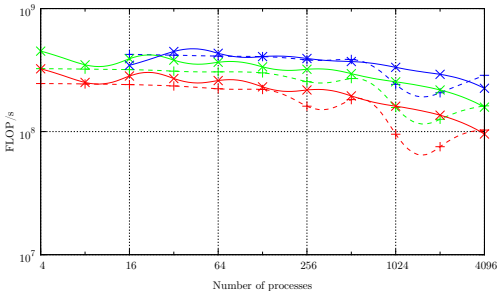


(b) Blue Gene

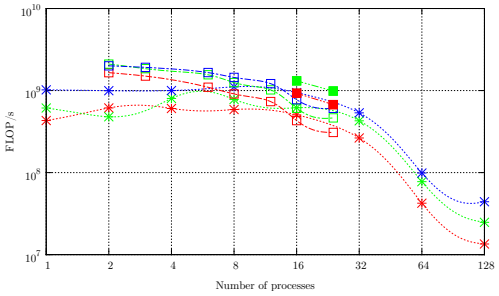


(c) Others

Figure 6: Results of the Dirac operator application test, plotting the average performance per MPI process. Symbols as on page 8.

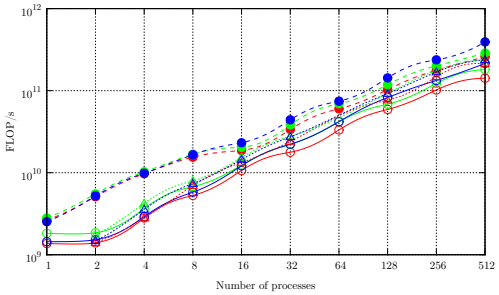


(a) HPC Wales and BlueIce2

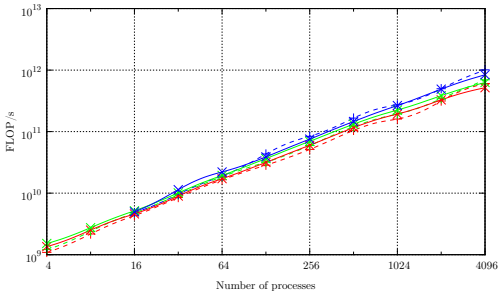


(b) Blue Gene

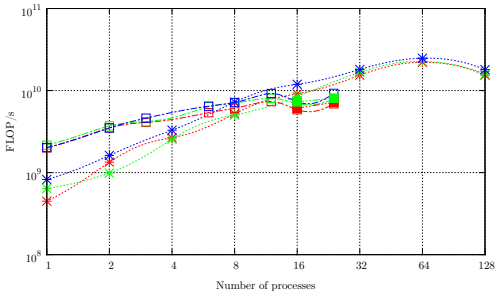


(c) Others

Figure 7: Total combined results of all three tests, plotted for the whole ensemble. Symbols as on page 8.

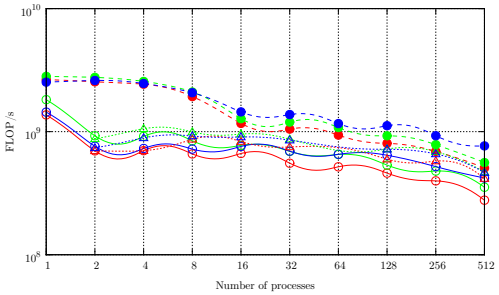


(a) HPC Wales and BlueIce2

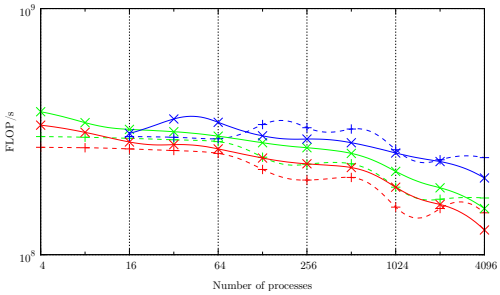


(b) Blue Gene

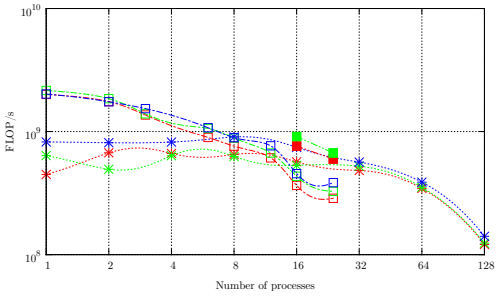


(c) Others

Figure 8: Total combined results of all three tests, plotting the average performance per MPI process. Symbols as on page 8.

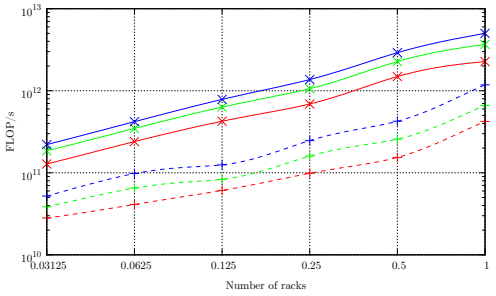


Figure 9: Rack-for-rack comparison between Blue Gene/P and Blue Gene/Q performance, using a larger lattice size in the case of the Blue Gene/Q to allow use of a full rack. Symbols as on page 8.

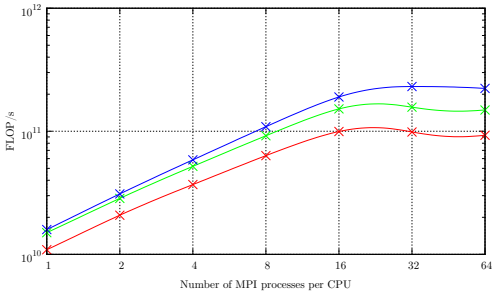


Figure 10: Performance comparison of a single Blue Gene/Q partition at various subdivisions; from one process per CPU through to 64 processes per CPU (four per core). Symbols as on page 8.

performance plateaus, showing little to no gain to using more than one thread per core.

It is worth drawing attention to some features of the data that are particularly interesting. For this discussion we will concentrate on the Dphi test, since it is most representative of "real-world" performance.

For the Blue Gene/P, unexpected upticks in performance were seen at 512 and 4096 cores (128 and 1024 nodes, or $\frac{1}{4}$ and 1 rack respectively) that were not seen on the Blue Gene/Q. This is most likely due to the lattice geometry aligning fortuitously with the 3D torus network—it would not be seen on Blue Gene/Q due to the different network topology and core density. No tuning of the process mapping was performed for the tests; the default alignments were used. More careful choice of the alignments may be able to give a similar (or better) speedup to other partition sizes.

The threading results on Blue Gene/Q are not unexpected. Since HiRep and BSMBench are single-threaded, subdividing the processors by core is equivalent to adding more cores to the problem. A significant speedup from multithreading is not expected: since multithreading only allows efficient context switching between threads rather than parallel execution (which is what multiple cores achieve), a large gain would only be possible if the code spent much of its time idle, which is not generally the case. A small gain can be found in a compute-intensive theory by having two processes rather than one per core; since this gain is effectively "free" there is no reason not to use it, unless the resulting local lattice is too small—in which case the extra overheads will outweigh the gains—or there are no more directions available to parallelise.

The Mac Pro system at 16 and 24 processes is also multi- (or Hyper-) threading; notice that the performance at 16 processes drops to approximately that of 8 processes. While the multiple threads per core can in principle execute concurrently (unlike on the Blue Gene/Q), in practise the shared FPU means that very little is parallelised, giving behaviour very close to what would be seen in a non-parallel multithreading core. The drop in performance, rather than a plateau, occurs for two reasons: firstly, the threads are localised to a single core; they cannot float between the cores and use whatever slack compute capacity there is. Secondly, the code has relatively frequent barriers to allow communication between the processes. Since the code makes heavy use of the floating-point units (which are shared between the "Hyper"-threads), these two factors have the effect of forcing the code to operate in lock-step on all cores, including those with only one task allocated, making the program run as slowly as if all cores had two threads running. Since the code has little idle time (even less than on the Blue Gene, since here there are no interconnect delays at all), no performance gains over 12 processes are seen at 24 processes. To test the lock-step theory, comms and balance tests[6] were performed with two 8- and 12-process runs in parallel with each other (for a total of 16 and 24 processes), and their FLOP counts added. This showed a slight enhancement over the single-run results, suggesting that the requirement to synchronise between processes does indeed reduce the potential gains from multithreading.

The HPC Wales machines, in particular the Sandy Bridge system, show an increasing separation between the three tests as the number of processes is increased. This nicely illustrates

[6]The compute test required more memory than was present in the Mac Pro, so the results were not of interest.

the assertion made earlier that the tests differ in their communications versus compute demands: the splitting begins once the processes no longer fit on the node, so must begin to use the interconnects, and increase as the interconnects are relied upon more. The slowdown is least severe for the compute, and most severe for the comms test, as we expect. This is not observed on the Blue Gene machines, whose architecture is designed from the ground up to be massively parallel, hence their advanced network topology and very high-speed interconnects. In the case of ULGQCD, the interconnects are too slow to make an informed judgement; however, the expected drop-off in performance is seen once the node size is exceeded.

The BlueIce2 system is very similar to the HPC Wales Westmere system in its makeup; their performance results are unsurprisingly also close to each other. The performance of BlueIce2 in the intermediary regime (from 16 cores where the job no longer fits on a single node, up to 256 cores) sits above the performance of the HPC Wales machine—this appears somewhat surprising, since its processors are slightly slower. In fact this enhancement is due to the job's parallelisation: on BlueIce2 it was possible to parallelise with 8 cores per node, and $N_{\text{proc}}/8 = 2^n$ nodes, whereas allocating a given number of cores on HPC Wales generally returns the cores on the minimum number of nodes. Since the lattice is parallelised in divisions of 2, keeping the powers of 2 in the parallelisation improves the communications performance significantly; this is notable in that the comms test is the one gaining the largest performance boost on BlueIce2. Confirming this explanation is the downtick in performance (to slightly below HPC Wales) at 512 cores: since BlueIce2 only has 50 nodes, the neat power-of-two parallelisation is no longer possible at the highest number of cores, meaning the communications speed enhancement is lost.

# 6 Conclusions

In this work we have presented BSMBench, a novel HPC benchmark suite derived from the latest research code used for Monte Carlo calculations in theoretical particle physics. BSMBench has the ability to vary the ratio between computation and communication demands. This property allows the tool to provide a robust assessment of the capabilities of different architectures under scenarios that can mimic a wide range of workloads. BSMBench can be downloaded from `http://www.bsmbench.org` and comes with an easy set of instructions and a limited set of dependencies in order to allow users not so familiar with parallel programming to deploy the suite on various multicore architectures. Although we have presented three particular realisations, other possibilities for obtaining a different ratio of communication over computation can be implemented changing a relatively small portion of the source code.

In addition to exploring other parallelisation strategies and porting to different architectures, which will be implemented following further developments of HiRep, the Lattice Gauge Theory code from which BSMBench is derived, other future directions for upcoming releases of the benchmark suite are being planned. Potential extensions to this project would be to implement a size-independent test, where the lattice size is chosen based on the machine size. Since the total number of FLOPs is proportional to the lattice volume, a reference figure could be given for a base $4^4$ lattice, and multiplied up to the machine size. Results compared to the current version would be approximately comparable. It would be beneficial to find a small set of quantities that would allow full characterisation of a machine's performance, to enable ranking of machines as in the TOP500 ranking (since comparison of graphs, while enlightening, is time-consuming). Finding out what classes of HPC use are characterised by which test would also expand the benchmark's utility. Finally, many benchmarks now come with the ability to analytically estimate the performance on a given machine from the machine's design characteristics; such a model for BSMBench would be a powerful tool for designing new systems, in particular once other codes' performance can be modelled in terms of BSMBench's tests.

# Acknowledgements

We thank Jim Sexton for invaluable insights in the early stage of this project. The simulations discussed in this work were performed on the STFC-funded DiRAC facilities, on High Performance Computing Wales machines and on systems supported by the Royal Society. EB acknowledges the support of an IBM student internship and of STFC through a Doctoral Training Grant. The work of BL is supported by STFC (grant ST/G000506/1) and by the Royal Society (grant UF09003). The research of CP is

supported by a Lundbeck Foundation Fellowship grant.

# A Lattice Gauge Theory

A complete review of the Lattice formalism, and the framework of Quantum Field Theory (QFT), which underlies it, can each span multiple textbooks, and thus are well outside the scope of this paper. Instead, this appendix will review the important results from the fields which could assist in understanding the remainder of this paper. Those wanting to learn more about QFT and Lattice techniques could do worse than looking in, for example, [24, 25] and [26, 27].

A mathematical topic that plays a central role in QFT is group theory. Once again, a full view of group theory is well beyond the scope of this paper, but for the interested reader there are many good textbooks introducing group theory in the context of physics [28, 29].

SU($N$) groups (which are the primary groups of interest in particle physics, and the only ones used in this paper) are defined as the groups of complex, unitary $N \times N$ matrices with determinant one. This defining representation (also called the fundamental representation) is not the only possibility to describe the abstract SU($N$) group, which may be represented also in terms of larger square matrices. Physically interesting higher representations include the symmetric, antisymmetric, and adjoint representations, which are called two-index representations, since each matrix in any of those representations can be expressed as a tensor product of two matrices in the fundamental representation. SU($N$) matrices (of any representation) can act on vectors (or conversely the vectors can be said to transform under a given representation of the group).

Each representation consists of $D_\mathrm{R} \times D_\mathrm{R}$ matrices, which act on $D_\mathrm{R}$-dimensional vectors, where $D_\mathrm{R}(N)$ is called the *dimension* of the representation R. We have already seen that for the fundamental representation, $D_\mathrm{F}(N) = N$. For the adjoint representation, $D_\mathrm{A}(N) = N^2 - 1$. For the anti-symmetric and symmetric representations, $D_\mathrm{AS}(N) = \frac{1}{2}N(N-1)$ and $D_\mathrm{S}(N) = \frac{1}{2}N(N+1)$ respectively.

Most observables calculable in QFT, and by extension in modern particle physics, may be expressed in the so-called path integral formulation, due to Feynman, as correlation functions in the form

$$\langle X \rangle = \frac{1}{Z} \prod_i \left( \int \mathrm{D}\xi_i \right) X \, e^{\mathrm{i}S[\xi_i]} \,, \qquad (1)$$

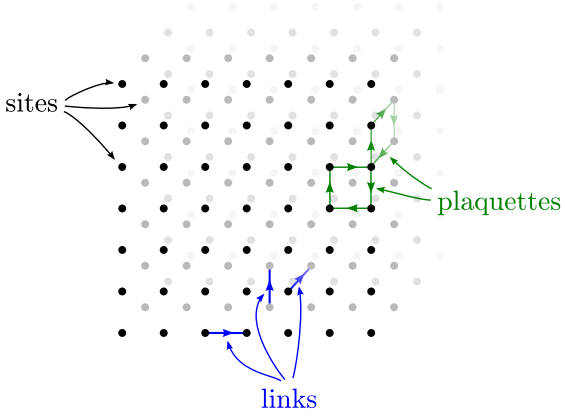


Figure 11: An example of a three-dimensional lattice, showing lattice sites, links between them, and elementary plaquettes formed from squares of links.

where $Z = \prod_i \left( \int \mathrm{D}\xi_i \right) e^{\mathrm{i}S}$ (the *path integral* or *partition function*) enforces $\langle 1 \rangle = 1$, $\prod_i \left( \int \mathrm{D}\xi_i \right)$ represents a functional integration over the space of fields $\{\xi_i\}$, i is the imaginary unit defined by $\mathrm{i}^2 = -1$, and $S[\xi_i]$ is the action defining the theory.

Most theories currently of interest contain gauge fields $A_\mu$ in the adjoint representation of the gauge group (most frequently chosen to be SU($N$)), and $N_\mathrm{f}$ fermion flavours $\psi_i$, spinor fields lying in some representation of the gauge group. Various choices for $N$ and $N_\mathrm{f}$ define different theories. At fixed $N$ and $N_\mathrm{f}$, the representation in which a field lies determines the physical properties of the corresponding system. For example, QCD has gauge group SU(3), and six fermion flavours in the fundamental representation. The same theory with the same number of adjoint quark fields would have vastly and noticeably different properties.

The functional integral present in Equation (1) represents an integration over all possible configurations of the fields, and is in general not analytically defined. We may, however, find a robust definition by discretising the problem—that is, placing it onto a finite grid, or *lattice*, with the functional integral of a field replaced by a product of integrals over the field values at each point; i.e.

$$\int \mathrm{D}\xi(x) \mapsto \prod_i \left( \int dx_{x_i} \right) . \qquad (2)$$

Figure 11 shows a three-dimensional example of such a lattice; however, the theories of interest in physics are most frequently four-dimensional, with three spatial directions and one temporal

direction, carrying an additional − sign in calculations of distance. Ultimately, it is because of this minus sign that the exponential in (1) is complex. Complex integrals are oscillating and give rise to cancellations that would be difficult to treat with numerical methods. We may simplify things somewhat by performing a *Wick rotation* into "imaginary time", where space and time are on the same footing, each making a positive contribution to distance calculations, just as the three spatial dimensions do in 3D calculations, leading to the setup being referred to as Euclidean spacetime. This has the effect of introducing a factor of i into the exponent in Equation (1). Now the exponential becomes damped, and for all practical purposes the integrals have support in a compact region in the integration variables. The original theory can be recovered by inverting the Wick rotation after the calculation has been performed. Although better behaved numerically after the Wick rotation, the set of integrals that needs to be computed rapidly becomes intractable using grid methods, since it scales linearly with the lattice volume (i.e. quartically with the lattice's extent), and so Monte Carlo techniques must be used.

For the lattice theories we are interested in, the Euclidean physical value of an observable $X$ can be expressed as

$$\langle X \rangle_{\text{Eucl.}} = \frac{1}{Z} \int \mathrm{D}U_\mu \mathrm{D}\phi_i \mathrm{D}\phi_i^\star \, X \, e^{-S_{\text{eff}}} \,, \quad (3)$$

with

$$S_{\text{eff}} = S_{\text{g}}[U_\mu] + \phi^\star \left( (D[U_\mu])^\dagger \, D[U_\mu] \right)^{-1} \phi \,. \quad (4)$$

Here $U_\mu$ (which represents the gauge variables and can be seen as the path-ordered exponentials of the $A_\mu$ along the link considered) is a field of SU($N$) matrices with support on the links of the lattice. $\phi$ is a vectorial quantity describing the fermions ($\phi^*$ being its complex conjugate) defined over all lattice sites. It transforms under the fundamental or a two-index representation of the gauge group (a flavour index running from one to $N_{\text{f}}$ identifying the fermion species is also understood). $D^\dagger$ is the Hermitian conjugate of the Dirac operator $D$ and an even number of flavours $N_{\text{f}}$ has been assumed throughout. Equations (3) and (4) expose the main computational kernel of the problem, which is the inversion of $H = (D[U_\mu])^\dagger \, D[U_\mu]$. In current calculations, this inversion dominates the calculation for the values of $N$ currently accessible in numerical simulations, with the part involving $S_{\text{g}}[U_\mu]$ only adds a modest computational overhead.

Together with the dimension of the group $N$, the representation in which the fermions transform determines the dimension of the Dirac operator, and hence the computational complexity of the problem. Given the sub-grid parallelisation strategy, these parameters also affect the size of the messages being transferred between neighbouring processors, and hence the communication requirements.

# References


[1] "HPL - a portable implementation of the High-Performance Linpack benchmark for distributed-memory computers." `http://www.netlib.org/benchmark/hpl`.

[2] "TOP500 supercomputer sites." `http://www.top500.org/`.

[3] J. Dongarra and M. A. Theroux, "Towards a new metric for ranking High Performance Computing systems," Tech. Rep. SAND2013-4744, Sandria National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California 94550, June, 2013. `http://www.sandia.gov/~maherou/docs/HPCG-Benchmark.pdf`.

[4] "The Graph 500 list." `http://www.graph500.org/`.

[5] "The Green500 list." `http://green500.org/`.

[6] **ATLAS** Collaboration, The ATLAS Collaboration, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," *Phys. Lett.* **B716** no. 1, (2012) 1 – 29, `arXiv:1207.7214`.

[7] **CMS** Collaboration, The CMS Collaboration, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC," *Phys. Lett.* **B716** no. 1, (2012) 30 – 61, `arXiv:1207.7235`.

[8] S. Schaefer, "Algorithms for lattice QCD: progress and challenges," *AIP Conf.Proc.* **1343** (2011) 93–98, `arXiv:1011.5641 [hep-ph]`.

[9] C. Pica, "Overview of iterative solvers for LQCD." `http://www.physik.uni-bielefeld.de/igs/schools/Tools2008/pica.pdf`.

[10] "DD-HMC algorithm for two-flavour lattice QCD." http://luscher.web.cern.ch/luscher/DD-HMC/index.html.

[11] "The Chroma library for lattice field theory." http://usqcd.jlab.org/usqcd-docs/chroma/.

[12] L. Del Debbio, A. Patella, and C. Pica, "Higher representations on the lattice: Numerical simulations. SU(2) with adjoint fermions," *Phys. Rev.* **D81** (2010) 094503, arXiv:0805.2058 [hep-lat].

[13] L. Del Debbio, B. Lucini, A. Patella, C. Pica, and A. Rago, "Conformal versus confining scenario in SU(2) with adjoint fermions," *Phys. Rev.* **D80** (2009) 074507, arXiv:0907.3896 [hep-lat].

[14] L. Del Debbio, B. Lucini, A. Patella, C. Pica, and A. Rago, "Infrared dynamics of Minimal Walking Technicolor," *Phys. Rev.* **D82** (Jul, 2010) 014510, arXiv:1004.3206.

[15] L. Del Debbio, B. Lucini, A. Patella, C. Pica, and A. Rago, "Mesonic spectroscopy of Minimal Walking Technicolor," *Phys.Rev.* **D82** (2010) 014509, arXiv:1004.3197 [hep-lat].

[16] S. Catterall, L. Del Debbio, J. Giedt, and L. Keegan, "MCRG Minimal Walking Technicolor," *PoS* **LATTICE2010** (2010) 057, arXiv:1010.5909 [hep-ph].

[17] B. Lucini, G. Moraitis, A. Patella, and A. Rago, "Orientifold Planar Equivalence: The Quenched Meson Spectrum," *PoS* **LATTICE2010** (2010) 063, arXiv:1010.6053 [hep-lat].

[18] F. Bursa, L. Del Debbio, D. Henty, E. Kerrane, B. Lucini, A. Patella, C. Pica, T. Pickup, and A. Rago, "Improved lattice spectroscopy of Minimal Walking Technicolor," *Phys. Rev.* **D84** (Aug, 2011) 034506, arXiv:1104.4301.

[19] R. Lewis, C. Pica, and F. Sannino, "Light asymmetric dark matter on the lattice: SU(2) technicolor with two fundamental flavors," *Phys. Rev.* **D85** (Jan, 2012) 014504, arXiv:1109.3513.

[20] E. Bennett and B. Lucini, "Topology of Minimal Walking Technicolor," *Eur. Phys. J.* **C73** no. 5, (2013) 1–7, arXiv:1209.5579. http://dx.doi.org/10.1140/epjc/s10052-013-2426-6.

[21] A. Hietanen, C. Pica, F. Sannino, and U. I. Sondergaard, "Orthogonal Technicolor with isotriplet dark matter on the lattice," *Phys.Rev.* **D87** no. 3, (2013) 034508, arXiv:1211.5021 [hep-lat].

[22] P. Vranas, G. Bhanot, M. Blumrich, D. Chen, A. Gara, P. Heidelberger, V. Salapura, and J. C. Sexton, "The Blue Gene/L supercomputer and Quantum ChromoDynamics, Gordon Bell finalist," in *Proceedings of SC06*. 2006. http://sc06.supercomputing.org/schedule/pdf/gb110.pdf.

[23] M. Lüscher, "Lattice QCD on PCs?," *Nucl. Phys. B Proc. Suppl.* **106–107** no. 0, (2002) 21 – 28, arXiv:hep-lat/0110007. LATTICE 2001 Proceedings of the XIXth International Symposium on Lattice Field Theory.

[24] M. E. Peskin and D. V. Schroeder, *An Introduction To Quantum Field Theory (Frontiers in Physics)*. Westview Press, 1995.

[25] A. Zee, *Quantum Field Theory in a nutshell*. Princeton University Press, 2010.

[26] H. J. Rothe, *Lattice Gauge Theories: An Introduction*. World Scientific, 2005.

[27] J. Smit, "Introduction to quantum fields on a lattice: A robust mate," *Cambridge Lect. Notes Phys.* **15** (2002) 1–271. http://www.slac.stanford.edu/spires/find/hep/www?j=00385,15,1.

[28] H. Jones, *Groups, representations, and physics*. Taylor & Francis, 1998.

[29] H. Georgi, *Lie algebras in particle physics: from isospin to unified theories*. Westview Press, 1999.